*Optimization Methods*

Draft of August 26, 2005

# II.
# Analysis of
# Linear Programming Solutions

**Robert Fourer**

***Department of Industrial Engineering and Management Sciences***
***Northwestern University***
***Evanston, Illinois 60208-3119, U.S.A.***

***(847) 491-3151***

4er@iems.northwestern.edu
http://www.iems.northwestern.edu/~4er/

## Introduction

By solving a linear program, you can learn more than just the values of the variables and of the objective function. As an introduction to some of the possibilities, chapter 8 first derives a fundamental "duality" property of linear programs. Chapter 9 then explains how you can interpret the values of certain "dual variables" to help you analyze the sensitivity of solutions to changes in costs, demands and other data.

# 8. Duality

The idea of duality is simple. Suppose we are given a certain formulation of linear program, such as the standard algebraic form of $m$ inequality constraints on variables $x_1, \ldots, x_{\mathcal{N}}$:

$$\begin{array}{lll} \text{Minimize} & \sum_{j=1}^{n} c_j x_j & \\ \text{Subject to} & \sum_{j=1}^{n} a_{ij} x_j \geq b_i, & i = 1, \ldots, m \\ & x_j \geq 0, & j = 1, \ldots, n \end{array}$$

Then there is a linear program related to this one, called its ***dual,*** having $n$ constraints on variables $\pi_1, \ldots, \pi_m$:

$$\begin{array}{lll} \text{Maximize} & \sum_{i=1}^{m} \pi_i b_i & \\ \text{Subject to} & \sum_{i=1}^{m} \pi_i a_{ij} \leq c_j, & j = 1, \ldots, n \\ & \pi_i \geq 0, & i = 1, \ldots, m \end{array}$$

If either one of these linear programs has a finite optimal value — so that it is neither unbounded nor infeasible — then the other one has exactly the same optimal value. Furthermore, if the original LP is solved by the simplex algorithm, then the vector $\pi$ computed at the last iteration turns out to be an optimal solution to the dual LP.

We begin below by defining and illustrating the duality relationship for different formulations of linear programs. Then we give simple proofs of duality and related properties for a few common forms.

## 8.1 Rules for forming the dual

The original linear program to be solved is customarily referred to as the ***primal*** problem. Thus above we gave a primal-dual pair, which could be written more compactly in matrix and vector terms as

| *Primal* | | *Dual* | |
|---|---|---|---|
| Minimize | $cx$ | Maximize | $\pi b$ |
| Subject to | $Ax \geq b$ | Subject to | $\pi A \leq c$ |
| | $x \geq 0$ | | $\pi \geq 0$ |

The dual of the dual linear program is again the primal. Thus there is a symmetry between the two problems; we choose to call one the primal only because it happens to be the formulation of original interest.

Given a different primal formulation, we get a different form of dual. For example, here is the dual for a primal in our standard algebraic form, with all equality constraints:

| *Primal* | | *Dual* | |
|---|---|---|---|
| Minimize | $cx$ | Maximize | $\pi b$ |
| Subject to | $Ax = b$ | Subject to | $\pi A \leq c$ |
| | $x \geq 0$ | | |

In the dual, the restrictions $\pi \geq 0$ have disappeared. We could present many different cases like this, but it will instead be most useful to give the general rules that allow a dual to be deduced from any primal.

In general terms, we can say that the primal has some number $m$ of constraints and some number $n$ of variables, and is representable as some kind of minimization:

$$\begin{array}{lll} \text{Minimize} & \sum_{j=1}^{n} c_j x_j & \\ \text{Subject to} & \sum_{j=1}^{n} a_{ij} x_j \,?\, b_i, & i = 1, \ldots, m \\ & x_j \,?\, 0, & j = 1, \ldots, n \end{array}$$

The dual, in turn, is representable as a maximization in $n$ constraints and $m$ variables:

$$\begin{array}{lll} \text{Maximize} & \sum_{i=1}^{m} \pi_i b_i & \\ \text{Subject to} & \sum_{i=1}^{m} \pi_i a_{ij} \,?\, c_j, & j = 1, \ldots, n \\ & \pi_i \,?\, 0, & i = 1, \ldots, m \end{array}$$

The question marks represent relationships that may vary from one variable to another and one constraint to another, but that are linked in two ways between the primal and the dual.

First, for each one of the constraints $\sum_{j=1}^{n} a_{ij} x_j \,?\, b_i$ in the primal LP, there is a variable $\pi_i$ in the dual. If the constraint is an equality, then $\pi_i$ is unrestricted in sign. If the constraint is a $\geq$ or $\leq$ inequality, then $\pi_i$ is restricted to be $\geq 0$ or $\leq 0$, respectively.

Second, for each one of the variables $x_j$ in the primal LP, there is a constraint $\sum_{i=1}^{m} \pi_i a_{ij} \,?\, c_j$ in the dual. If $x_j$ is restricted to be $\geq 0$ (as is the case in most of our models) then the corresponding dual constraint is $\sum_{i=1}^{m} \pi_i a_{ij} \leq c_j$. If $x_j$ is unrestricted in sign, then the dual constraint becomes an equality, while if $x_j$ is restricted to be $\leq 0$ then the dual constraint is a $\geq$ inequality.

All of these observations can be summarized together in a table of duality rules:

| *Primal* | *Dual* |
|---|---|
| Minimize $\sum_{j=1}^{n} c_j x_j$ | Maximize $\sum_{i=1}^{m} \pi_i b_i$ |
| $\sum_{j=1}^{n} a_{ij} x_j \geq b_i$ | $\pi_i \geq 0$ |
| $\sum_{j=1}^{n} a_{ij} x_j = b_i$ | $\pi_i$ unrestricted |
| $\sum_{j=1}^{n} a_{ij} x_j \leq b_i$ | $\pi_i \leq 0$ |
| $x_j \geq 0$ | $\sum_{i=1}^{m} \pi_i a_{ij} \leq c_j$ |
| $x_j$ unrestricted | $\sum_{i=1}^{m} \pi_i a_{ij} = c_j$ |
| $x_j \leq 0$ | $\sum_{i=1}^{m} \pi_i a_{ij} \geq c_j$ |

If instead the primal is a maximization, then the table looks like this:

| *Primal* | *Dual* |
|---|---|
| Maximize $\sum_{j=1}^{n} c_j x_j$ | Minimize $\sum_{i=1}^{m} \pi_i b_i$ |
| $\sum_{j=1}^{n} a_{ij} x_j \le b_i$ | $\pi_i \ge 0$ |
| $\sum_{j=1}^{n} a_{ij} x_j = b_i$ | $\pi_i$ unrestricted |
| $\sum_{j=1}^{n} a_{ij} x_j \ge b_i$ | $\pi_i \le 0$ |
| $x_j \ge 0$ | $\sum_{i=1}^{m} \pi_i a_{ij} \ge c_j$ |
| $x_j$ unrestricted | $\sum_{i=1}^{m} \pi_i a_{ij} = c_j$ |
| $x_j \le 0$ | $\sum_{i=1}^{m} \pi_i a_{ij} \le c_j$ |

The dual is now a minimization, and the ≤ and ≥ signs exchange places in the constraint entries of the table.

## 8.2 Examples

As illustrations of particular duality rules, we use one small linear program made up for the purpose, and one from a game theory application that we previously developed. A linear programming model for the transportation problem is then used to show how a class of primal models gives rise to a certain class of dual models.

*A small linear program.* As an initial example, consider the following specific primal LP:

$$
\begin{array}{ll}
\text{Minimize} & 1x_1 + 2x_2 + 3x_3 + 4x_4 + 5x_5 \\
\text{Subject to} & 11x_1 + 12x_2 + 13x_3 + 14x_4 + 15x_5 \ge 101 \\
& 21x_1 + 22x_2 + 23x_3 + 24x_4 + 25x_5 \le 202 \\
& 31x_1 + 32x_2 + 33x_3 + 34x_4 + 35x_5 = 303 \\
& x_1 \ge 0,\ x_2 \ge 0
\end{array}
$$

The coefficients have been artificially chosen to highlight the relationship to the dual. When the above rules are applied, the dual linear program looks like this:

$$
\begin{array}{ll}
\text{Maximize} & 101\pi_1 + 202\pi_2 + 303\pi_3 \\
\text{Subject to} & 11\pi_1 + 21\pi_2 + 31\pi_3 \le 1 \\
& 12\pi_1 + 22\pi_2 + 32\pi_3 \le 2 \\
& 13\pi_1 + 23\pi_2 + 33\pi_3 = 3 \\
& 14\pi_1 + 24\pi_2 + 34\pi_3 = 4 \\
& 15\pi_1 + 25\pi_2 + 35\pi_3 = 5 \\
& \pi_1 \ge 0,\ \pi_2 \le 0
\end{array}
$$

Whereas the primal is a minimization with three constraints in five variables, the dual is a maximization with five constraints in three variables.

The dual variables $\pi_1$, $\pi_2$ and $\pi_3$ correspond to the three primal constraints. Since the first two primal constraints are ≥ and ≤ inequalities, the first two dual variables are restricted to be nonnegative and nonpositive, respectively. The third primal constraint is an equality, however, so $\pi_3$ is unrestricted in sign.

The "right-hand side" constants 101, 202 and 303 of the primal constraints become the corresponding objective function coefficients of the dual variables.

Each of the five dual constraints corresponds to one of the five primal variables. Since the first two primal variables are restricted to be nonnegative, the first two dual constraints are ≤ inequalities; since the remaining primal variables are unrestricted in sign, the last three dual constraints are equalities. The coefficients in each dual constraint are taken exactly from the constraint coefficients of the corresponding primal variable, while the right-hand side of each dual constraint is the objective coefficient of the primal variable. For example, the numbers in the fourth dual constraint $14\pi_1 + 24\pi_2 + 34\pi_3 = 4$ come from the terms $14x_4$, $24x_4$ and $34x_4$ in the primal constraints, and the term $4x_4$ in the primal objective.

Overall, you can think of the dual as being the primal turned "on its side". The dual's coefficient matrix is just the transpose of the primal's (recall that $\pi A$ is the same as $A^T\pi$). Every column in the primal becomes a row in the dual, and every row in the primal becomes a column in the dual. Also, the objective function coefficients in the primal become the right-hand sides in the dual, while the right-hand side constants in the primal becomes the objective coefficients in the dual.

*A linear program from game theory.* In *Optimization Models* we showed how certain desirable strategies for two-person zero-sum games could be derived as the solutions to linear programs. In the particular game used as an example, player $\mathcal{A}$'s random strategy was given by the following LP:

$$\begin{array}{ll}
\text{Maximize} & z \\
\text{Subject to} & z \le 4x_1 - x_2 - x_3 \\
& z \le -2x_1 + 4x_2 - 2x_3 \\
& z \le -3x_1 - 3x_2 + 4x_3 \\
& x_1 + x_2 + x_3 = 1 \\
& x_1 \ge 0,\ x_2 \ge 0,\ x_3 \ge 0
\end{array}$$

We want to show that the corresponding dual is precisely the linear program for player $\mathcal{B}$'s random strategy. Then we will have an explanation for the fact, which we previously observed, that the two linear programs have the same optimal value.

To apply our table of duality rules, we rearrange the constraints so that all variables appear to the left:

$$\begin{array}{ll}
\text{Maximize} & 0x_1 + 0x_2 + 0x_3 + 1z \\
\text{Subject to} & -4x_1 + 1x_2 + 1x_3 + 1z \le 0 \\
& 2x_1 - 4x_2 + 2x_3 + 1z \le 0 \\
& 3x_1 + 3x_2 - 4x_3 + 1z \le 0 \\
& 1x_1 + 1x_2 + 1x_3 + 0z = 1 \\
& x_1 \ge 0,\ x_2 \ge 0,\ x_3 \ge 0
\end{array}$$

There are four constraints in the four variables $(x_1, x_2, x_3, z)$. Thus the dual will also have four constraints in four variables. If we call the dual variables

$(y_1, y_2, y_3, w)$, then the duality rules for a maximization give the following linear program:

$$\begin{array}{ll} \text{Minimize} & 0y_1 + 0y_2 + 0y_3 + 1w \\ \text{Subject to} & -4y_1 + 2y_2 + 3y_3 + 1w \geq 0 \\ & 1y_1 - 4y_2 + 3y_3 + 1w \geq 0 \\ & 1y_1 + 2y_2 - 4y_3 + 1w \geq 0 \\ & 1y_1 + 1y_2 + 1y_3 + 0w = 1 \\ & y_1 \geq 0,\ y_2 \geq 0,\ y_3 \geq 0 \end{array}$$

The coefficients in the objective come from the right-hand side values in the primal. The numbers in the first three constraints come from the coefficients of $x_1$, $x_2$ and $x_3$, while the numbers in the fourth constraint come from the coefficients of $z$.

It is easy to rearrange this dual so that it is in a more appealing form, like the original primal:

$$\begin{array}{ll} \text{Minimize} & w \\ \text{Subject to} & w \geq 4y_1 - 2y_2 - 3y_3 \\ & w \geq -y_1 + 4y_2 - 3y_3 \\ & w \geq -y_1 - 2y_2 + 4y_3 \\ & y_1 + y_2 + y_3 = 1 \\ & y_1 \geq 0,\ y_2 \geq 0,\ y_3 \geq 0 \end{array}$$

This is exactly the linear program that we devised to compute player $\mathcal{B}$'s random strategy. Since the two players' LPs are duals, they must give the same optimal objective value; thus we found that the max-min expected winnings of $\mathcal{A}$ (equal to $z$) were the same as the min-max expected losses of $\mathcal{B}$ (equal to $w$).

It can be shown that this duality is not a property just of the particular game that we analyzed. The players' linear programs are always duals, for any two-person zero-sum game, and hence their expected payoffs are always equal.

*A transportation model.* As the game theory example suggests, it is useful to be able to state the dual of a general linear programming model, rather than for just a particular linear program. Here we use a transportation model as an example.

As seen in *Optimization Models,* the data for a simple transportation model can be written in terms of an origin set $\mathcal{O}$ and a destination set $\mathcal{D}$ as

- $a_o$ units of material available at origin $o$, for each $o \in \mathcal{O}$
- $b_d$ units of material required at destination $d$, for each $d \in \mathcal{D}$
- $c_{od}$ cost per unit to ship material from origin $o$ to destination $d$, for each $o \in \mathcal{O}$ and $d \in \mathcal{D}$

The decision variables are

- $x_{od}$ units of material that are shipped from origin $o$ to destination $d$, for each $o \in \mathcal{O}$ and $d \in \mathcal{D}$

The formulation as a linear program is as follows:

$$\begin{array}{lll} \text{Minimize} & \sum_{o\in\mathcal{O}}\sum_{d\in\mathcal{D}} c_{ij}x_{ij} & \\ \text{Subject to} & \sum_{d\in\mathcal{D}} x_{od} \le a_o, & \text{for all } o \in \mathcal{O} \\ & \sum_{o\in\mathcal{O}} x_{od} \ge b_d, & \text{for all } d \in \mathcal{D} \\ & x_{od} \ge 0, & \text{for all } o \in \mathcal{O} \text{ and } d \in \mathcal{D} \end{array}$$

This is a general statement of the primal linear programming model for the transportation problem. What is the corresponding general form of the dual?

To begin, observe that the primal has an origin or "supply' constraint for each member of the set $\mathcal{O}$, and a destination or "demand" constraint for each member of the set $\mathcal{D}$. There will have to be a dual variable associated with each of these constraints:

$\sigma_o$ dual variable associated with the $i$th supply constraint, for each $o \in \mathcal{O}$

$\lambda_d$ dual variable associated with the $j$th demand constraint, for each $d \in \mathcal{D}$

Since the supply constraints are $\le$ inequalities, while the demand constraints are $\ge$ inequalities, the rules for dualizing a minimization say that all $\sigma_i \le 0$, while all $\lambda_j \ge 0$.

The dual objective function — to be maximized — must sum the products of the dual variables times their corresponding right-hand-side values. Hence the formula for the dual objective is $\sum_{o\in\mathcal{O}} \sigma_o a_o + \sum_{d\in\mathcal{D}} \lambda_d b_d$.

There must be a dual constraint corresponding to each primal variable $x_{od}$. Since the primal has $x_{od} \ge 0$, the form of the corresponding constraint is

$$\pi \cdot (\text{constraint coefficients of } x_{od}) \le (\text{objective coefficient of } x_{od}).$$

The constraint coefficients of $x_{od}$ consist of the coefficients in the supply and in the demand constraints, and the objective coefficient is $c_{od}$, so the dual constraint becomes:

$$\sum_{o\in\mathcal{O}} \sigma_o \cdot (\text{coeff of } x_{od} \text{ in supply constraint for origin } o) +$$
$$\sum_{d\in\mathcal{D}} \lambda_d \cdot (\text{coeff of } x_{od} \text{ in demand constraint for destination } d) \le c_{od}.$$

But the coefficient of $x_{od}$ in all the supply constraints is zero, except for the supply constraint at origin $o$, in which the coefficient is one. Similarly, the coefficient of $x_{od}$ in all the demand constraints is zero, except for a one in the demand constraint for destination $d$. Hence the above expression collapses to $\sigma_o \cdot 1 + \lambda_d \cdot 1 \le c_{od}$.

Putting the objective and constraints together, we have the following dual model:

$$\begin{array}{lll} \text{Maximize} & \sum_{o\in\mathcal{O}} \sigma_o a_o + \sum_{d\in\mathcal{D}} \lambda_d b_d & \\ \text{Subject to} & \sigma_o + \lambda_d \le c_{od}, & \text{for all } o \in \mathcal{O} \text{ and } d \in \mathcal{D} \\ & \sigma_o \le 0, & \text{for all } o \in \mathcal{O} \\ & \lambda_d \ge 0, & \text{for all } d \in \mathcal{D} \end{array}$$

This linear program will give the same objective value as the primal, for any choice of $\mathcal{O}$, $\mathcal{D}$, $a_o$, $b_d$, and $c_{od}$. If there are $m$ origins and $n$ destinations, then whereas the primal has $m+n$ constraints in $mn$ variables, the dual has $mn$ constraints in only $m+n$ variables. We will later come back to this dual to illustrate some interpretations of dual variables.

## 8.3 Duality theorems

We now proceed to show why the primal and dual linear programs really do always have equal objective values. In the process, we derive several other useful results pertaining to the optimal dual values.

To keep things simple, we will use the following form of primal-dual pair:

| *Primal* | | *Dual* | |
|---|---|---|---|
| Minimize | $cx$ | Maximize | $\pi b$ |
| Subject to | $Ax = b$ | Subject to | $\pi A \le c$ |
| | $x \ge 0$ | | |

Much the same results apply to other common formulations, however.

*Weak duality.* It is easy to show that the objective value of the primal LP can never be pushed to less than the value of the dual LP, or vice versa:

> If $\bar{x}$ is feasible for the primal, and $\bar{\pi}$ is feasible for the dual, then $c\bar{x} \ge \bar{\pi}b$.

This is proved by observing that, for any feasible $\bar{x}$ and $\bar{\pi}$,

$$A\bar{x} = b \quad \Rightarrow \quad \bar{\pi}A\bar{x} = \bar{\pi}b$$

$$\bar{\pi}A \le c \text{ and } \bar{x} \ge 0 \quad \Rightarrow \quad \bar{\pi}A\bar{x} \le c\bar{x}$$

Putting together the two expressions deduced at the right, we get the desired result.

Since optimal solutions must be feasible, we have an immediate corollary:

> If $x^*$ is optimal for the primal, and $\pi^*$ is optimal for the dual, then $cx^* \ge \pi^* b$.

This is a step in the right direction, which is why it is called "weak" duality. What we eventually want to show is that $cx^* > \pi^* b$ can never happen, so that the primal and dual objectives must always satisfy $cx^* = \pi^* b$.

Before we go further, however, there are several other useful corollaries that are easily derived. Suppose that we discover feasible solutions $\bar{x}$ and $\bar{\pi}$ such that $c\bar{x}$ does equal $\bar{\pi}b$. Then no other feasible solution $\hat{x}$ can give a lower primal objective value; if it did, then we would have $c\hat{x} < \bar{\pi}b$, violating weak duality. By the same reasoning, no other feasible solution $\hat{\pi}$ can give a higher dual objective value. It follows that both $\bar{x}$ and $\bar{\pi}$ must be optimal:

> If $\bar{x}$ is feasible for the primal and $\bar{\pi}$ is feasible for the dual, and if $c\bar{x} = \bar{\pi}b$, then $\bar{x}$ is optimal for the primal and $\bar{\pi}$ is optimal for the dual.

It turns out that the simplex algorithm finds exactly an $\bar{x}$ and a $\bar{\pi}$ of this kind, which will be the key to proving strong duality.

What if the primal LP is unbounded? Then we can find a feasible solution $\bar{x}$ that makes $c\bar{x}$ as small as we like. In particular, if $\bar{\pi}$ were any feasible solution for the dual, then we could find an $\bar{x}$ so that $c\bar{x} < \bar{\pi}b$, violating weak duality. Hence:

> If the primal is unbounded, then the dual has no feasible solution.

Analogous reasoning gives the opposite result:

> If the dual is unbounded, then the primal has no feasible solution.

The converses of these statements do not necessarily hold, however. It is possible for both a primal linear program and its dual to have no feasible solutions.

*Strong duality.* Suppose now that we are given a primal linear program in the standard form, and that it has a finite optimal value. Then we have seen that the simplex algorithm can be applied, and that it must eventually stop in step (2). Let $\hat{x}$ denote the basic solution when the algorithm stops, and let $\hat{\pi}$ denote the vector that was computed in step (1).

Certainly $\hat{x}$ is feasible for the primal LP, because the simplex algorithm only generates feasible solutions. What about $\hat{\pi}$? From step (1) we have $\hat{\pi}B = c_{\mathcal{B}}$, which means that

$$\hat{\pi} \cdot a_i = c_i \text{ for each variable } x_i \text{ that is basic.}$$

Because the algorithm has stopped in step (2), we must also have all reduced costs $d_j = c_j - \hat{\pi}a_j \geq 0$, which rearranges to give

$$\hat{\pi} \cdot a_j \leq c_j \text{ for each variable } x_j \text{ that is nonbasic.}$$

Putting these together, we have that $\hat{\pi}a_j \leq c_j$ for *every* variable $x_j$, or equivalently $\hat{\pi}A \leq c$, so that $\hat{\pi}$ turns out to be feasible for the dual LP.

Finally, what about the objective values given by $\hat{x}$ and $\hat{\pi}$? Since $\hat{x}$ is the current basic solution, it is given by $B\hat{x}_{\mathcal{B}} = b$, $\hat{x}_{\mathcal{N}} = 0$. Combining this with $\hat{\pi}B = c_{\mathcal{B}}$, we see that

$$c\hat{x} = c_{\mathcal{B}}\hat{x}_{\mathcal{B}} = (\hat{\pi}B)\hat{x}_{\mathcal{B}} = \hat{\pi}(B\hat{x}_{\mathcal{B}}) = \hat{\pi}b.$$

The objective values are equal.

We have shown that the simplex algorithm stops with a primal feasible solution $\hat{x}$ and a dual feasible solution $\hat{\pi}$, which give equal objective values. Thus, by the corollary to weak duality cited above, both of these solutions are optimal — and, as a result, the dual has achieved the same objective value as the primal. We have thus demonstrated the following strong duality property:

> If the primal has a finite optimal value, then the dual has the same optimal value.

The opposite, exchanging primal and dual, can be shown by similar means.

You might have noticed that, in the above argument, we used only the fact that the simplex algorithm eventually stops in step (2). Hence this can be used as an alternative proof that the simplex algorithm has an optimal solution when it stops.

A further interesting consequence of duality concerns the solving of systems of equations and inequalities. If you can find $x^*$ and $\pi^*$ that together solve

$$\begin{aligned} &Ax^* = b,\ x^* \geq 0 \\ &\pi^* A \leq c \\ &cx^* = \pi^* b \end{aligned}$$

then weak duality implies that you will have simultaneous optimal solutions $x^*$ for the primal and $\pi^*$ for the dual. Furthermore, strong duality implies that, if the primal or the dual has a finite optimal solution, then you can *always* find the solutions to both by just solving the above inequalities. In other words, if you have any method for solving inequalities, you can use that method to solve linear programs as well. It can be concluded that solving inequalities is *just as hard* as solving linear programs — which justifies using the simplex algorithm to find feasible solutions in phase I as well as optimal solutions in phase II, as previously explained.

## 8.4 Complementary slackness

One more alternative condition for optimality is also useful in analyzing some solutions. For motivation, imagine that the simplex method has stopped in step (2), and has thus found optimal solutions $x^*$ and $\pi^*$.

If $x_j^* > 0$, then $x_j$ must be a basic variable at the optimal solution. Hence $a_j$ is one of the columns in the basis matrix $B$, and the equations $\pi^* B = c_{\mathcal{B}}$ imply that $\pi^* a_j = c_j$. We can conclude that $x_j^* > 0$ implies that there is no slack in the $j$th dual constraint. Or, equivalently, if there is slack in the $j$th dual constraint, then $x_j^* = 0$.

This kind of conclusion is called a ***complementary slackness*** result. It can be extended to the following conditions that hold for all optimal solutions:

> Feasible solutions $x^*$ for the primal and $\pi^*$ for the dual are optimal if and only if:
>
> > Either $x_j^* = 0$ or there is no slack in the $j$th dual constraint (or both), for every $j = 1, \ldots, n$.
> >
> > Either $\pi_i^* = 0$ or there is no slack in the $i$th primal constraint (or both), for every $i = 1, \ldots, m$.

If the primal form has $Ax = b$, then there never is any slack in the primal constraints, and the second set of conditions above holds trivially. But if the primal has $Ax \geq b$, then there can be slack in the primal as well as dual constraints.

To see why this property holds, suppose that the primal and dual are minimization and maximization problems subject to inequality constraints, and suppose further that $x^*$ and $\pi^*$ are feasible solutions. Then feasibility can be expressed as

$$\begin{array}{llll} a^i x^* \geq b_i, & i = 1, \ldots, m & \pi^* a_j \leq c_j, & j = 1, \ldots, n \\ x_j^* \geq 0, & j = 1, \ldots, n & \pi_i^* \geq 0, & i = 1, \ldots, m \end{array}$$

Now, starting from the strong duality result and rearranging, we have

$$\begin{array}{lll} x^*, \pi^* \text{ optimal} & \Leftrightarrow & cx^* = \pi^* b \\ & \Leftrightarrow & cx^* - \pi^* A x^* = \pi^* b - \pi^* A x^* \\ & \Leftrightarrow & (c - \pi^* A) x^* = \pi^* (b - A x^*) \\ & \Leftrightarrow & \sum_{j=1}^n (c_j - \pi^* a_j) x_j^* + \sum_{i=1}^m \pi_i^* (a^i x^* - b_i) = 0 \end{array}$$

In the last line, primal and dual feasibility insure that each term $(c_j - \pi^* a_j) x_j^* \geq 0$, and each term $\pi_i^* (a^i x^* - b_i) \geq 0$. Thus we have a sum of nonnegative terms being zero, which can only be so if each individual term is zero. That is:

$$\begin{array}{lll} x^*, \pi^* \text{ optimal} & \Leftrightarrow & (c_j - \pi^* a_j) x_j^* = 0 \text{ for each } j = 1, \ldots, n, \text{ and} \\ & & \pi_i^* (a^i x^* - b_i) = 0 \text{ for each } i = 1, \ldots, m \end{array}$$

These products can be zero, however, if and only if at least one of the factors in each is zero. Thus the condition becomes that either $x_j^* = 0$ or $c_j = \pi^* a_j$, which means no slack in the $j$th dual constraint; and either $\pi_i^* = 0$ or $a^i x^* = b_i$, which means no slack in the $i$th primal constraint.

It can happen that both $x_j^* = 0$ *and* there is no slack in the $j$th dual constraint. Thus we cannot conclude that a zero value for $x_j^*$ implies a positive value for the slack, or that a zero value for the slack implies a positive value for $x_j^*$. However, it turns out that there is always at least one solution for which such a conclusion is true:

> If the primal and dual have finite optimal values, then there exist optimal solutions $x^*$ and $\pi^*$ such that:
>
> > For every $j = 1, \ldots, n$, either $x_j^* = 0$ and there is slack in the $j$th dual constraint, or $x_j^* \neq 0$ and there is no slack in the $j$th dual constraint.
> >
> > For every $i = 1, \ldots, m$, either $\pi_i^* = 0$ and there is slack in the $i$th primal constraint, or $\pi_i^* \neq 0$ and there is no slack in the $i$th primal constraint.

This is known as the *strict* complementary slackness property.

# 9. Interpretation of Dual Prices

We begin this topic somewhat indirectly, by showing how the primal values provide information about the sensitivity of an optimal solution to changes in the primal objective coefficients. Then we show how the dual values provide analogous information about sensitivity to changes in the primal right-hand-side values—which are also the dual objective coefficients. This analysis leads to some other interpretations in terms of the "value" of certain materials or resources.

## 9.1 Sensitivity to objective function coefficients

Suppose that you have already found an optimal solution $x^*$ and an optimal value $z^*$ for some linear program. What do you ask next? One common concern is the "sensitivity" of the optimum to various changes in the data. Here we consider, in particular, how the optimal value might change as some objective function coefficient $c_d$ changes.

It is easy to imagine that, if $c_d$ does not change too much, then the optimal basis $B$ stays the same. As a result, the optimal solution given by $Bx^*_{\mathcal{B}} = b$ is unchanged — since it depends only on the basis $B$ and on $b$, but not on $c_d$ at all. In particular, then, if $c_d$ is changed to $c'_d = c_d + \delta_d$, then the optimal value changes to

$$z' = c'x^* = c_1x^*_1 + \cdots + (c_d + \delta_d)x^*_d + \cdots + c_nx^*_n = z^* + \delta_d x^*_d.$$

The rate of change in the objective is given by $x^*_d$. Or, to put it another way,

> As long as the current basis stays optimal, the objective value changes by $x^*_d$ for every unit change in $c_d$.

In calculus terms, you can think of $x^*_d$ as equaling the derivative of $z^*$ with respect to $c_d$.

If $c_d$ is either increased or decreased by a sufficiently large amount, then the same basis will no longer be optimal. The optimal solution will be different, and so will the rate of change in the objective. Nevertheless, we can imagine plotting the optimal value against different values of $c_d$. The following is typical:

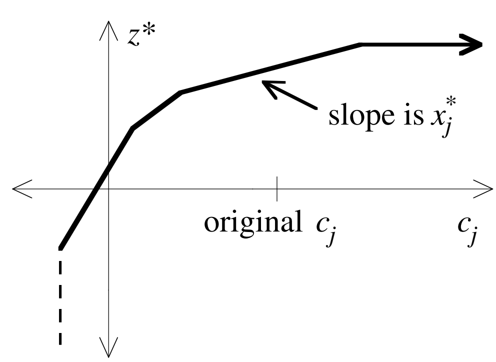

For any minimizing linear program, this graph turns out to be *concave* and *piecewise-linear* as above. (For a maximization it would be the same but convex.) In the neighborhood of the original $c_d$, the graph is linear with slope $x_d^*$, for the reasons explained above. But as $c_d$ gets larger, the slope of the graph tends to get smaller; while as $c_d$ gets smaller, the slope tends to get larger. At some point, as at the left above, the slope may become "infinite", so that the graph effectively stops there.

To understand the implications of this graph, it is useful to interpret it in terms of a simple model like the one that arises from the diet problem:

$$\begin{array}{ll} \text{Minimize} & cx \\ \text{Subject to} & Ax \geq b \\ & x \geq 0 \end{array}$$

As usual, $c_d$ represents the cost per unit of food $j$, $b_o$ is the units required of nutrient $i$, and $a_{ij}$ is the units of nutrient $i$ in each unit of food $j$. The optimal value $x_d^*$ is the number of units of food $j$ purchased.

The graph shows that, at first, the total cost of the diet goes up by $x_d^*$ cents for every cent increase in the cost $c_d$ of the $j$th food. However, as $c_d$ increases further, the cost of the diet goes up at a slower rate. In fact, for sufficiently large values of $c_d$ — where the graph is flat at the right — the cost of the diet shows no change for any further increase. This is a situation you would expect; if the $j$th food becomes costly enough, then it is not used at all in the diet ($x_d$ is nonbasic) and subsequent increases in its cost do not affect the objective value.

The graph also shows that, at first, the total cost of the diet goes down by $x_d^*$ cents for every cent decrease in the cost $c_d$ of the $j$th food. However, as $c_d$ decreases further, the cost of the diet goes down at a faster rate. For a sufficiently small value of $c_d$ (negative in this case) the objective value drops suddenly to $-\infty$; this just means that the linear program becomes unbounded.

## 9.2 Sensitivity to right-hand-side values

Now consider what happens if some constant $b_o$ is changed. Even if the basis $B$ stays optimal, the solution to $Bx_{\mathcal{B}} = b$ will change, and hence the effect on the objective value is not immediately evident.

Here we need to look at the dual linear program. Since the optimal dual solution $\pi^*$ satisfies $\pi^* B = c_{\mathcal{B}}$, a change in $b_o$ will not affect it. So if $b_o$ is changed to $b_o' = b_o + \delta_o$, the optimal value of the dual LP changes to

$$z' = \pi^* b' = \pi_1^* b_1 + \cdots + \pi_o^* (b_o + \delta_o) + \cdots + \pi_m^* b_m = z^* + \pi_o^* \delta_o.$$

By the strong duality property, this also represents the change in the optimal value of the primal LP. The rate of change in the objective is given by $\pi_o^*$. Or, to put it another way,

> As long as the current basis stays optimal, the objective value changes by $\pi_o^*$ for every unit change in $b_o$.

In calculus terms, you can think of $\pi_o^*$ as equaling the derivative of $z^*$ with respect to $b_o$.

We can now imagine a second graph, in which the optimal value is plotted against values of $b_o$:

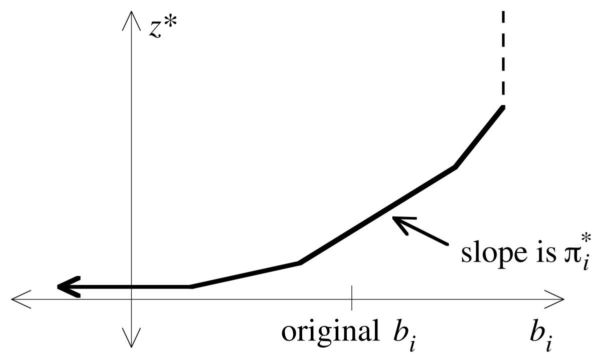


For any minimizing linear program, this graph turns out to be *convex* and piecewise-linear as above, the opposite of the case for $c_d$. (For a maximization it would be the same but concave.) In the neighborhood of the original $b_o$, the graph is linear with slope $\pi_o^*$, for the reasons explained above. But as $b_o$ gets larger, the slope of the graph tends to get larger; while as $b_o$ gets smaller, the slope tends to get smaller. At some point, as at the right above, the graph may also stop with an “infinite” slope.

Using the diet problem again as an example, a change in $b_o$ represents a change in the requirement for the $i$th nutrient. At first, the total cost of the diet goes down by $\pi_o^*$ cents for every unit decrease in the requirement. However, as $b_o$ decreases further, the cost of the diet goes down at a slower rate. In fact, for sufficiently small values of $b_o$ — where the graph is flat at the left — the cost of the diet shows no change for any further decrease. This is again a situation you would expect; if the requirement for the $i$th nutrient becomes small enough, then it is optimal to supply more of that nutrient than is needed, and further decreases in the requirement have no effect.

The graph also shows that, at first, the total cost of the diet goes up by $\pi_o^*$ cents for every unit increase in the requirement for the $i$th nutrient. However, as the requirement increases further, the cost of the diet goes up at a faster rate. For a sufficiently large value of $b_o$, the objective value jumps suddenly to $+\infty$; this signifies that the linear program has no feasible solution when too much of the nutrient is required. (Actually there is always a feasible solution when the only constraints are $Ax \geq b$ and nonnegativity of the variables; but infeasibility can occur when, for example, there are upper bounds on the variables.)

## 9.3 Interpretations of the optimal dual values

We have already shown one interpretation of the optimal solution $\pi^*$ to the dual linear program. Each $\pi_o^*$ gives the sensitivity of the optimal objective to changes in a right-hand-side value $b_o$. The situation can be summarized as follows:

When the primal is a minimization:

> Each increase of 1 in $b_o$ causes an increase of at least $\pi_o^*$ in the optimal objective value.
>
> Each decrease of 1 in $b_o$ causes a decrease of at most $\pi_o^*$ in the optimal objective value.

When the primal is a maximization:

> Each increase of 1 in $b_o$ causes an increase of at most $\pi_o^*$ in the optimal objective value.
>
> Each decrease of 1 in $b_o$ causes a decrease of at least $\pi_o^*$ in the optimal objective value.

This interpretation can be applied to any right-hand side of any linear program. It should be kept in mind that, if $\pi_o^* < 0$, then an "increase" of $\pi_o^*$ is in fact a decrease in the amount of $-\pi_o^*$.

A related way of interpreting $\pi^*$ can be illustrated with the diet problem. Suppose that someone offered to sell you some units of nutrient $i$. For each unit you bought, you would be able to reduce the requirement $b_o$ by one — since you would need one less unit of nutrient to be provided by the foods. For each unit that $b_o$ was reduced, however, the optimal cost of the foods would go down by at most $\pi_o^*$. Thus, it stands to reason, you would be willing to pay at most $\pi_o^*$ for each extra unit of nutrient $i$.

Under this interpretation, $\pi_o^*$ is the ***marginal value*** of the $i$th nutrient. The dual solution can similarly be interpreted in terms of values for many other kinds of linear programs. For instance, in a maximum-revenue production problem, $\pi_o^*$ may be the marginal value of the $i$th scarce resource; it is the most that the objective will increase for each unit increase in the amount $b_o$ of the resource available, and hence the most you would be willing to pay to get more of the resource. In summary, the possible interpretations are as follows:

When the primal is a minimization:

> If $\pi_o^* \geq 0$, then $\pi_o^*$ is the most you would be willing to pay to decrease $b_o$ by one unit.
>
> If $\pi_o^* \leq 0$, then $-\pi_o^*$ is the most you would be willing to pay to increase $b_o$ by one unit.

When the primal is a maximization:

> If $\pi_o^* \geq 0$, then $\pi_o^*$ is the most you would be willing to pay to increase $b_o$ by one unit.
>
> If $\pi_o^* \leq 0$, then $-\pi_o^*$ is the most you would be willing to pay to decrease $b_o$ by one unit.

The particular significance of increasing or decreasing $b_o$ by one unit will depend on the particular model that you are analyzing.

The transportation model, together with its dual introduced in the previous chapter, provide a good example of how optimal dual values can be interpreted in a particular situation. Each dual variable $\sigma_o$ provides information about sensitivity of the optimum to $a_o$, the amount of material available at origin $o$; while each dual variable $\lambda_d$ provides information about sensitivity to $b_d$, the amount of material required at destination $d$.

Using the general principles derived above, we can make the following specific statements about the relationship between $a_o$ and the optimal value $\sigma_o^*$:

- Each increase of 1 in $a_o$ causes an increase of at least $\sigma_o^*$ in the optimal objective value.
- Each decrease of 1 in $a_o$ causes a decrease of at most $\sigma_o^*$ in the optimal objective value.

These are somewhat misleading, however, because in fact the constraints of the dual LP require that $\sigma_o^* \leq 0$. An "increase of at least $\sigma_o^*$" is really a *decrease* of *at most* an amount $-\sigma_o^* \geq 0$, and similarly for a decrease. So we can restate the interpretation of $\sigma_o^*$ as follows:

- Each increase of 1 in $a_o$ causes a decrease of at most $-\sigma_o^* \geq 0$ in the optimal objective value.
- Each decrease of 1 in $a_o$ causes an increase of at least $-\sigma_o^* \geq 0$ in the optimal objective value.

This corresponds to what you would expect. Increasing the amount of material $a_o$ available at origin $o$ can only give you more shipping options, thereby causing a decrease in total shipment cost. On the other hand, decreasing the amount of material available has the opposite effect, causing an increase in total cost.

There is a similar relationship between $b_d$ and the optimal dual value $\lambda_d^*$:

- Each increase of 1 in $b_d$ causes an increase of at least $\lambda_d^*$ in the optimal objective value.
- Each decrease of 1 in $b_d$ causes a decrease of at most $\lambda_d^*$ in the optimal objective value.

In this case the dual constraints require $\lambda_d^* \geq 0$, so there is no further adjustment to make. Intuitively, an increase in the amount of material $b_d$ required at destination $d$ will tend to cause an increase in total shipping cost, while a decrease in the amount required will tend to cause a decrease in shipping cost.

Now consider how you would react to offers allowing you to buy additional units of material. If you could buy additional units at origin $o$, then you could increase $a_o$ by that much; the optimal value would be reduced by up to $-\sigma_o^*$ per unit bought. On the other hand, if you could buy additional units at destination $d$, then you would be able to decrease $b_d$ by that much—because you would not need to ship as much to $d$ any more—and the objective would be decreased by up to $\lambda_d^*$ per unit bought. These are the marginal values of the material:

- You would be willing to pay at most $-\sigma_o^*$ per unit to buy extra material at origin $o$.
- You would be willing to pay at most $\lambda_d^*$ per unit to buy extra material at destination $d$.

Another of the constraints of the dual LP requires that $\sigma_o^* + \lambda_d^* \le c_{od}$. With a little rearranging, this says that $\lambda_d^* \le -\sigma_o^* + c_{od}$: the value of material at destination $d$ cannot exceed the value at origin $o$ by more than the shipment cost between them.

An even stronger statement can be made, by using the complementary slackness property. Since $x_{od}$ is the variable that corresponds to the dual constraint $\lambda_d \le -\sigma_o + c_{od}$, we have that either $x_{od}^* = 0$ or $\lambda_d^* = -\sigma_o^* + c_{od}$, or both. In other words,

- If $x_{od}^* > 0$, then $\lambda_d^* = -\sigma_o^* + c_{od}$.
- If $\lambda_d^* < -\sigma_o^* + c_{od}$, then $x_{od}^* = 0$.

This says that if material is actually shipped from origin $o$ to destination $d$ in some optimal solution, then its value at the destination is *exactly* its value at the origin plus the shipping cost. Hence, also, if the value of material at some destination $d$ is strictly less than the value at some origin $o$ plus the shipping cost, then it cannot be optimal to ship any material from $o$ to $d$.

Complementary slackness also applies to the dual variables $\sigma_o$ and their corresponding primal supply constraints. Either $\sigma_o^* = 0$ or $\sum_{d \in \mathcal{D}} x_{od}^* = a_o$, or both. Hence,

- If $\sigma_o^* < 0$, then $\sum_{d \in \mathcal{D}} x_{od}^* = a_o$.
- If $\sum_{d \in \mathcal{D}} x_{od}^* < a_o$, then $\sigma_o^* = 0$.

This says that if there is a positive marginal value $-\sigma_o^*$ to material at origin $o$, then an optimal solution will ship out all material currently available. Equivalently, if not all material is shipped out, then there is no value to obtaining more of it at origin $o$.